

Amendments to the Specification

Please amend paragraph [0002] on pg. 1 with the following amended paragraph:

[0002] A storage service provider may maintain a large network, such as a Fibre Channel Storage Area Network (SAN), to service the computing needs for one or more customers. The SAN includes numerous host systems including the customer applications linked via a Fibre Channel fabric to one or more storage systems, such as [[or]] one or more interconnected disk drives configured as a Redundant Array of Independent Disks (RAID), Just a Bunch of Disks (JBOD), Direct Access Storage Device (DASD), etc. Typically, a customer will pursue a service level agreement (SLA) with the storage service provider concerning the criteria under which network storage resources are provided, such as the storage capacity, network throughput, I/O response time, I/O operations per second, and other performance criteria under which the network resources will be provided. In certain situations, multiple customers with different levels of requirements specified in their service level agreements will share the same network resources. This requires that the storage service provider monitor and manage the network resources to ensure that the different customer requirements specified in the different service level agreements are satisfied.

Please amend para. [0009] on pg. 3 of the Specification with the following amended paragraph:

[0009] Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIGs. 1 and 2 illustrate network computing environments in which embodiments of the invention are implemented;

FIGs. 3, 4, 5, 6, 7, and 8 illustrate an arrangement of information on I/O paths between hosts and storage volumes and service level requirements providing performance criteria for service level agreements and associations therebetween in accordance with implementations of the invention;

FIGs. 9, 10, 11, and 12 illustrate operations performed to utilize the information described with respect to FIGs. 3-8 to manage the network shown in FIGs. 1 and 2; and

FIG. 13 illustrates a computing architecture that may be used to implement the network components described with respect to FIGs. 1 and 2.

FIG. 14a illustrates how the service level agreement elements are defined in an XML format.

FIG. 14b illustrates how the service level commitment information is implemented in the XML format.

FIG. 14c illustrates how the service level guarantee (SLG) information is implemented in the XML format.

Please amend the paragraph at lines 3-11 of the Specification with the following amended paragraph.

[0018] The storage service provider may further define, using the SLA client 28, a plurality of service level guarantees that define performance guarantees defining the level of performance the storage ~~service~~ service provider must provide pursuant to one or more service level agreements. A defined service level guarantee may apply to one or more application storage groups to define the level of service and performance expected of the connections identified by the application service connections (ASCs) included in the application service groups (ASGs) to which the service level guarantee is assigned. FIG. 5 illustrates information included in a service level guarantee record 90, including:

Please amend the paragraphs at lines 4-9 on pg. 9 of the Specification with the following amended paragraph:

~~Retention~~ Collection Interval 130: the time period during which the system will gather statistic data of ASGs defined by value and unit.

Reporting Interval 132: The interval at which the system will send collected statistics on performance data of the ASGs, defined by value and unit, to the performance analyzer 22 in the SLA server ~~[[18]]~~ 16. The collection and reporting intervals 130 and 132 may be adjusted.

Please amend paragraph [0021] on pg. 10 of the Specification as follows:

[0021] The storage service provider may review a service level agreement for a customer and then assign, using the SLA client 28, service level commitments to application service groups of application service connections for a customer by defining service level commitments for that customer. The storage service provider may enter information on connections (ASCs), groups of connections (ASGs), performance criteria (SLGs), and the relation therebetween (SLCs) at the SLA client 28, where the defined ASCs, ASGs, SLGs, and SLCs are stored in the SLA database ~~[[18]]~~ 20. Each instance of the above records (e.g., ASC, ASG, SLG, and SLC records) may be implemented in an Extensible Markup Language (XML) file or records within a database.

Please amend paragraph [0025] on pg. 10 of the Specification as follows:

[0025] FIG. 9 illustrates operations performed by the SLA server 16 when receiving (at block 200) performance data gathered by one performance gateway 14b, the SLA server 16 locates (at block 202) the performance data for the ASC 172a, 172b...172n (FIG. 9) for which the performance data was received. The located performance data for the ASC 172a, 172b...172c is updated (at block 204) with the new received performance data. In this way, the performance data may maintain each instance of a measured performance parameter, such as response time for each I/O request, I/O operations per second, I/O throughput.

Please amend paragraph [0031] on pg. 14 of the Specification as follows:

[0031] If (from the no branch of block 286) there are under performing ASCs, but no over performing ASCs to throttle, then the SLA server 16 generates (at block 290) an alert to notify the storage service provider of the underperforming ASCs. The storage service provider may be notified through the ~~admin~~ SLA client ~~[[18]]~~ 28.

Please amend paragraph [0032] on pg. 14 of the Specification as follows:

[0032] FIG. 13 illustrates an implementation of the SLA server 16 and SLA client ~~[[18]]~~ 28 in a web service based architecture. The SLA client 300 may include a web browser 302 that renders an SLA administrative graphical user interface (GUI) 304

through which the network administrator may interact with the SLA server 320. The SLA client 300 further includes a performance monitor 306 component that presents the response time in real-time to the network administrator from the performance information gathered by the SLA server 320. The SLA server 320 includes an HTTP server 322 to enable communication with the client web browser 302 and a Java servlet 324 to transfer information received at the HTTP server 322 to an application server 326 and to transfer information from the application server 326 through the HTTP server 322 to the client web browser 302. For instance, the application server 326 may transfer real-time performance data from the SLA database 328 to the client performance monitor 306